

Calibration of the Improved Stratospheric and Mesospheric Sounder. Part II: Optical Alignment, Spatial, and Pressure Modulator Calibrations

C. W. P. PALMER AND P. VENTERS

Department of Physics, University of Oxford, Oxford, United Kingdom

R. J. KNIGHT, J. BALLARD, AND T. J. NIGHTINGALE

Atmospheric Sciences Division, Rutherford Appleton Laboratory, Chilton, Oxfordshire, United Kingdom

P. E. MORRIS

Department of Physics, University of Oxford, Oxford, United Kingdom

(Manuscript received 22 April 1996, in final form 20 June 1996)

ABSTRACT

In Part I the Improved Stratospheric and Mesospheric Sounder instrument and its calibration facility were described, and the results of the radiometric and spectral calibrations were presented. In Part II the remaining prelaunch calibrations are discussed.

The optical calibrations include the procedure used to coalign the different spectral channels and the mapping of the resulting instantaneous fields of view. These maps revealed detailed stray light processes within the instrument. The remaining optical calibration is of angular motion of the instrument scan mirror. The pressure modulator calibrations consist of the filling procedures and algorithm validation, which allow the gas conditions within the modulator to be determined from the instrument telemetry. Some overall conclusions to the calibration process are drawn.

1. Introduction

This is the second of two papers describing the prelaunch calibration of the Improved Stratospheric and Mesospheric Sounder (ISAMS), a multichannel gas-correlation radiometer on the *Upper Atmosphere Research Satellite*. The first paper (Ballard et al. 1996) gave a detailed description of the instrument and the results of the radiometric and spectral calibrations. The present paper describes measurements relating to the instantaneous field of view (IFOV) and pointing of the instrument, and the calibration of the pressure-modulated cells (PMC), which implement the gas correlation. Initial scientific results and data validation are given in a series of recent papers (Connor et al. 1996; Dudhia and Livesey 1996; Goss-Custard et al. 1996; Lambert et al. 1996; Lopez-Valverde et al. 1996; Reburn et al. 1996; Remedios et al. 1996; Rodgers et al. 1996; Smith et al. 1996).

Only the optical system of the instrument will be

described here. Infrared radiation emitted by the earth's atmosphere viewed at the limb is collected by a telescope. The field image of the telescope is divided vertically into two subfields by a wedged mirror, and each subfield divided spectrally into four to produce eight final image planes, numbered from 0 to 7. In each of these, the radiation is detected by four-element arrays, to produce the nominal composite IFOV shown in Fig. 1. The individual elements of each channel have a rectangular IFOV with a 7:1 aspect ratio ($21' \times 3'$, equivalent to about $16.4 \text{ km} \times 2.3 \text{ km}$ in the atmosphere), separated by $3'$, so that the composite IFOV of the whole channel is square. The separation of the two subfields is also $3'$, with channels 0–3 nominally coaligned and viewing the upper subfield, and channels 4–7 nominally coaligned and viewing the lower. The optical train between the field-splitting mirror and each final image plane includes a PMC containing the target gas of the channel, as shown in Table 1. In channels 3 and 7 there is also a four-position filter wheel, giving four distinct spectral channels distinguished by a second digit, running from 0 to 3. The PMC gas in these channels is CO_2 , but the target gases for some of the filters are gases that cannot be contained in a PMC, such as O_3 .

The calibration was carried out in two campaigns in 1989 and 1990. The calibration facility at the University

Corresponding author address: Dr. C. W. P. Palmer, Department of Physics, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom.
E-mail: palmer@atm.ox.ac.uk

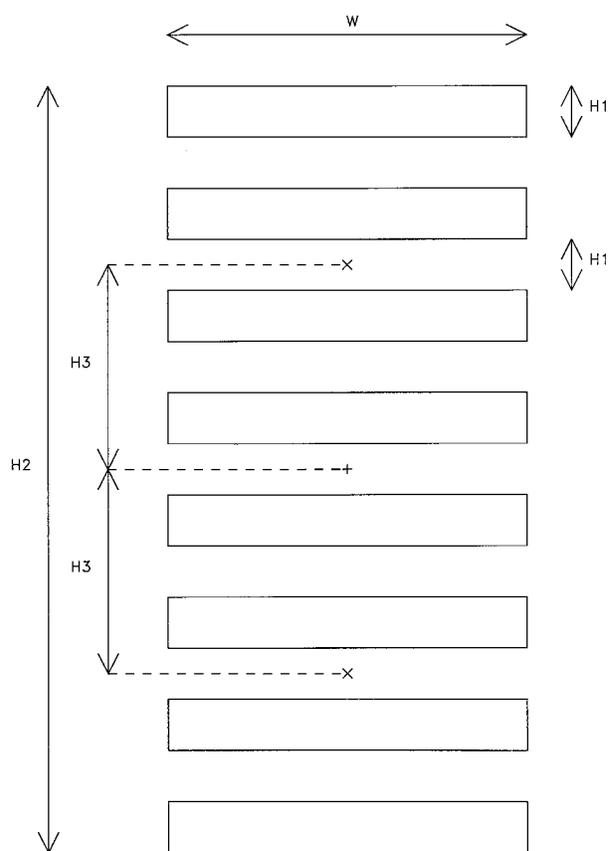


FIG. 1. Nominal IFOV: $H1 = 3'$, $H2 = 45'$, $H3 = 12'$, $W = 21'$.

of Oxford is described in Ballard et al. (1996). The optical measurements described in the next two sections were made using the collimator part of the collimator-monochromator system (CMS), which delivered a full-aperture beam from a scanning collimator to the instrument. All test equipment control and data acquisition was done using the ISAMS microprocessor system (IMS) and a modified form of the flight data system, so that all timing was controlled by a single clock. The common timing unit was the ISAMS measurement period (IMP, 1 IMP = 2.048 s) in which one sample was transmitted from each data channel (some of which, in the case of engineering data, were multiplexed).

The optical calibrations determine the geolocation and resolution of each radiance sample. These are of much greater significance for a limb-viewing instrument than for a nadir-viewing instrument, because the samples are always being taken from a region at the earth's limb with a very large radiance gradient. The next two sections describe measurements made on the alignment and IFOV of the optical system, section 4 describes the calibration of the angle encoders on the scan mirror, and section 5 the calibration of the PMCs. The final section draws together some conclusions from the two calibration papers.

TABLE 1. Channel assignments and 50% spectral widths.

| Channel | Target gas | 50% points | (cm^{-1}) |
|---------|-------------------------------|------------|----------------------|
| 0 | CO | 2176 | 2257 |
| 1 | H ₂ O | 1444 | 1515 |
| 2 | N ₂ O | 1262 | 1298 |
| 30 | CO ₂ | 628 | 655 |
| 31 | CO ₂ | 605 | 624 |
| 32 | HNO ₃ | 860 | 898 |
| 33 | O ₃ | 990 | 1010 |
| 4 | NO | 1856 | 1932 |
| 5 | NO ₂ | 1595 | 1629 |
| 6 | CH ₄ | 1320 | 1372 |
| 70 | CO ₂ | 628 | 655 |
| 71 | CO ₂ | 605 | 625 |
| 72 | N ₂ O ₅ | 1224 | 1249 |
| 73 | Aerosol | 818 | 833 |

2. Alignment

The four-element detector arrays are mounted into detector packages, which also include a folding mirror and two lenses. Translation and rotation of this package on its mounting surface, and rotation of the folding mirror, allowed correction for various alignment errors channel by channel; namely, defocus, image rotation, and incorrect image alignment. Measurement and correction of misalignment formed a significant part of the first calibration campaign.

The collimator was used with a 50- μm pinhole (equivalent to about 12") as the source aperture, which generates a diffraction- and aberration-limited beam. This was scanned in elevation across the center line of the IFOV, in a smooth scan taking about 45 min. This scan was repeated with the instrument scan mirror stepped $\pm 7.5'$ in azimuth, which moves the line of the scan toward the edges of the detectors. As the collimator beam scans across the IFOV, responses are seen from the channels in turn. When an element is not being illuminated by the beam, it is viewing the faceplate of the collimator, so that the main response induced by the collimator beam is observed on a background that varies with the temperature structure of the faceplate. In the longer wavelength channels, this background structure could be a significant fraction of the size of the wanted signal, because although the difference between the global and faceplate temperatures far exceeded the temperature differences between different parts of the faceplate, the global only filled a tiny fraction of the IFOV of an element. To remove the background structure, a slow chopper was implemented by a shutter mounted across the pinhole, which was opened and closed on alternate IMPs by a program running in the IMS.

Analysis of alignment data consisted of extraction of the "pinhole" signal, by subtracting "chopper closed" signals from the adjacent "chopper open" signals, and locating the edges of the main response, defined as the points with maximum positive and negative gradient of signal with respect to elevation angle. In this way, the

upper and lower edges of each detector element field image were located at three different points along each edge. Averaging these elevations produced an elevation for the channel boresight (marked “×” in Fig. 1). The datum for correct boresight location was defined on the first alignment map by a mask inserted at the field image plane just before the field-splitting mirror and centered on the telescope boresight (marked “+” in Fig. 1). Any elevation trend across the field indicated an image rotation, while the defocus was measured by repeating the central alignment scan twice with the collimator defocused in both senses by equal amounts. The value of the maximum signal gradient varied with defocus, and the sharpest image should occur at infinity focus. In practice, the best focus position varied across the image because of field curvature and small tilts between the plane of the detectors and the field plane, and an average value of the defocus was used. It turned out to be possible to measure the state of alignment much more accurately than it was possible to implement the small movements to the detector packages to correct any errors, and it took several iterations of this process before all detectors were within tolerance.

3. Field-of-view mapping

a. Measurements

The alignment scans revealed a lot of detailed structure in the field images, but it was not possible to build up an overall picture of the IFOV in this way, for a number of reasons. First, because the collimator moved continuously, the signal and background measurement did not quite coincide, so that a measurement artifact proportional to the spatial and/or temporal derivative of the background remained. Second, while for alignment purposes we could use the filter wheel positions giving largest collimator signal—typically the shortest wavelength filters 33 and 72—there was still a requirement to map the IFOV of the other filters, and in order to do this for the 15- and 16- μm CO_2 channels (30, 31, 70, and 71) it was necessary to increase the signal-to-noise ratio. Finally, the resolution of the alignment scans was not required for the overall IFOV map and would have meant an unnecessarily long measurement time.

For these reasons a different procedure was adopted for the IFOV map. The 50- μm diameter pinhole was replaced by a 50 μm \times 200 μm rectangular aperture, accurately aligned with its long dimension parallel to the collimator azimuth scan direction by viewing the collimator output beam with a theodolite. This immediately gave larger signals in all channels and reduced the resolution in the azimuth direction. A step-and-settle scan pattern was adopted so that signal and background measurements were collocated, and several measurements at each step location were taken to reduce the noise. The map was built up by repeating stepped scans in elevation at a range of azimuths to cover the whole

IFOV. Ideally, the step sizes should have been such that the images of the source aperture on successive elevation and azimuth steps just touched, but step sizes about 50% larger than this were used to reduce the time taken. Even so the step sizes were about one-tenth of the width of the nominal IFOV in elevation and about one-eighteenth in azimuth, and comparable with the smallest expected aberration/diffraction spot sizes, so it is reasonable to assume that no significant structure has been missed.

The problem of the low signal-to-noise ratio in the long wavelength channels remained. The signal in the short wavelength channels, particularly channels 0 and 4, increases exponentially with increasing source temperature, while in the long wavelength channels it only increases linearly, so that the shortwave channels overran before adequate signal was available in the longwave channels. This problem was handled by repeating the entire map twice at two different global temperatures. Map A consisted of 30 azimuth scans of 213 elevation steps with a source temperature of about 560 K. The filter wheels were set to 33 and 72 throughout, and each step lasted 8 IMPs. The source shutter opened and closed on alternate IMPs, as for the alignment scans, and the first IMP after a step (shutter open) was discarded to allow time for settling of the collimator structure. A time-symmetric difference of signal minus background was used, by giving half weight to the first and last background measurements on each step, which removes slow temperature changes of the collimator faceplate to first order. This map took 32 h.

Map B used a much higher source temperature of around 1000 K to increase the signals in the long wavelength channels and a longer time per step of 16 IMPs. In addition, the map coverage was extended slightly, with 31 scans, and increased coverage at the low elevation end, so that the total time taken increased to 75 h. At each step, one signal IMP was discarded for settling time as before, and the remaining 15 divided between the three remaining filter wheel positions of channels 3 and 7. This allowed time-symmetric differencing in one case, but not in the other two. The overranging of the other channels, particularly channels 0 and 4, was so severe that it was necessary to disconnect the channel 0, 1, 4, and 5 detector packages to prevent the excess current drawn by their preamplifiers affecting the gain of the channel 3 and 7 detectors. Channels 2 and 6, where the problem was not so severe, were left connected in order to verify the coregistration of the two maps.

Long-term gain changes were corrected by viewing the internal blackbody calibration source during the fly-back periods between elevation scans. The intrinsic stability of the source (a global with a regulated power supply) is perfectly adequate, but the instrument gain deteriorated over several days because of slow warming of the detectors. These blackbody signals were found to be a linear function of indicated detector temperature

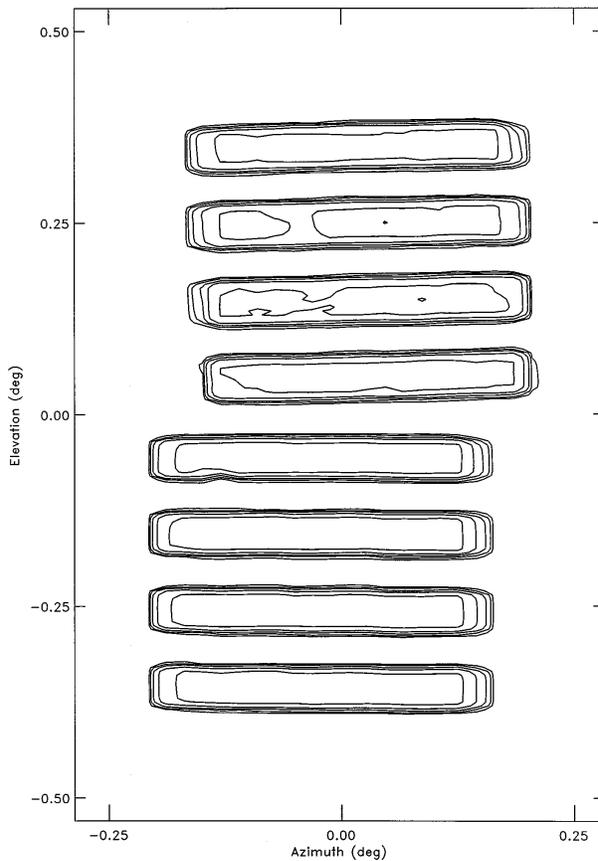


FIG. 2. Two-dimensional IFOV map for channels 2 and 4.

in each channel over the period of the maps, which therefore allowed the indicated temperature to be used as a proxy for a gain measurement to normalize out the long-term gain variation within, as well as between, scans.

The signals at each step, extracted as described above, were located on a measurement grid. Alignment scans, as described in section 2, were also carried out, enabling the telescope boresight to be located on this grid, and the rotation between the two elevation axes to be measured. The data were then relocated relative to the boresight and instrument axes, interpolated and regridded on a grid uniform in elevation, with $0.3'$ spacing. This dataset forms the primary IFOV database; because the retrieval process is directed mainly at retrieving vertical structure, the IFOV function is integrated over azimuth to yield one-dimensional IFOV functions and finally normalized to unity summed over all grid points. The incorporation of these azimuth-averaged IFOV shapes into the retrieval software is described by Dudhia and Livesey (1996) and Rodgers et al. (1996).

b. Discussion of IFOV results: Stray light

The two-dimensional IFOV measurements have extremely high signal-to-noise ratios, between about 500

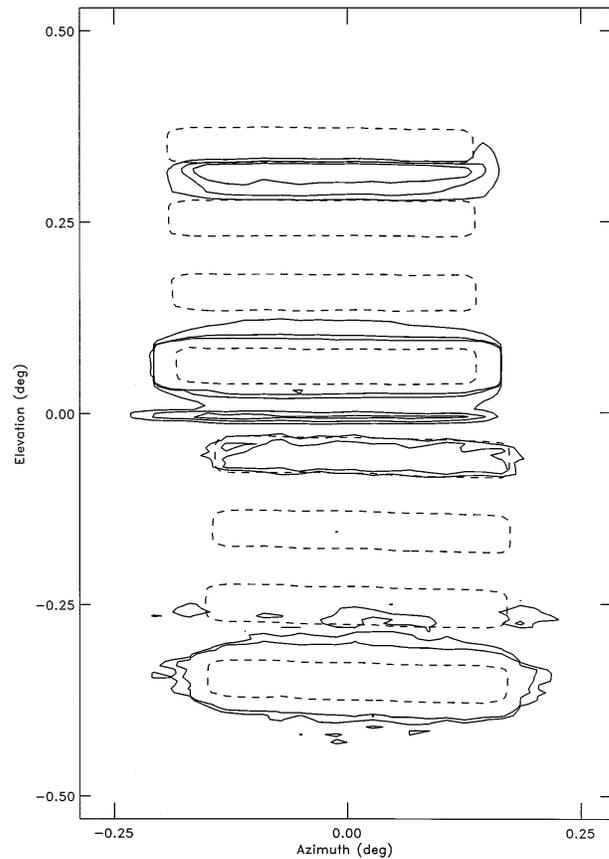


FIG. 3. Low-level responses in the IFOV map in channels 0 and 72 (see text for details).

for the long wavelength channels and 5×10^4 for the shortest wavelength channels. They thus reveal in great detail many aspects of the optical system. At the lowest level of detail, each detector element is a tightly focused rectangle, illustrated in Fig. 2 for channels 2 and 4, with contours at 5%, 10%, 20%, 40%, and 80% of the peak response per pixel. The azimuth offset is a feature of channels 0 and 4: These channels have InSb detector arrays, procured from a different manufacturer from the other arrays, and the array is slightly offset on the substrate. The residual uncorrected image rotation in channel 2 is also apparent in this figure.

At a higher level of detail, stray light effects associated with multiple reflections between the surface of the detector array and the final lens are observed. Two examples are shown in Fig. 3. In the upper subfield are plotted the 50% contours for the elements of channel 0, and the 0.5%, 1%, and 1.5% contours for element 3. In addition to the expected region around the element, two additional disconnected regions of low-level response are apparent: a region between elements 0 and 1, and a narrow line of response below element 3. Regions of increased response between detector elements are observed in most channels, and result from radiation that is reflected from the substrate surface and

from the final lens in the detector package. Line features are also observed in several channels, sometimes associated with the edges of other detector elements, but particularly in channels 0 and 4 with a line outside the region of the array altogether. Some of these out-of-field responses were initially quite significant, and an attempt was made to mask out the strongest ones in an accessible field image plane. This was very successful in channel 0 and moderately successful in channel 4, where the residual line feature has an integrated response of 1%.

In the lower subfield in Fig. 3 are plotted the 50% contours for the four elements of channel 72, and the 1% and 1.5% contours for element 3, which exhibit weak cross-talk, with response correlating with element 0. This feature is not present in the IFOV for the other filter wheel settings, so it is clearly optical in origin; since this is the shortest wavelength filter wheel channel (8.1 μm), the effect is almost certainly due to loss of efficiency in the antireflection coating over the detectors, since the coating is optimized for the 15- μm temperature-sounding channels.

At the highest level of detail some response is observed all over the subfield in which a detector is located, at least in the shorter wavelength channels, but this falls by over an order of magnitude as zero elevation is crossed, where the two subfields are divided by a wedged mirror at the final image plane of the telescope. In general, the responses in the opposite subfield are consistent with zero, with two exceptions. First, the elements nearest to zero elevation exhibit fringing responses that cross the field-splitting line, and the most natural interpretation of these is that a small degree of scattering around the specular direction has taken place at the telescope mirrors, which are the only place in the optical system where such small angle effects would be detectable in the final image. Finally, the background level in the "zero" subfield is typically not quite consistent with zero, but with an offset of order a quarter of the single sample noise level. Similar effects can also be seen if a spurious signal is extracted from the internal blackbody calibration views, using the time differencing appropriate to the IFOV measurements, and these offsets are clearly measurement artifacts. A truncated IFOV function was derived by setting elevation ranges for each element of each channel, within which they were nonzero, and renormalizing the truncated function. These ranges were chosen to include all the real low-level responses attributable to multiple reflections and scattering but to eliminate the regions consistent with zero.

4. Pointing

a. Requirements

The scan mirror mechanism has been briefly described in Part I. Its function is to scan the instrument line of sight (LOS) through the atmosphere to measure

radiance profiles from about the tropopause to a zero space view at about 150 km; in addition to direct the LOS either forward or backward of the $\pm Y$ axis to introduce sufficient spacecraft velocity along the LOS to cancel the velocity of the earth's rotation at the tangent point. The mirror is driven by stepper motors and lead screws in both axes.

The action of the plane scan mirror on the IFOV can be succinctly stated, but the consequences are quite subtle. If the IFOV function is measured when the scan mirror unit normal vector is \mathbf{m}_0 , and the mirror is subsequently rotated so that its normal is \mathbf{m} , then the IFOV function is rigidly rotated about an axis parallel to $\mathbf{m}_0 \times \mathbf{m}$; if the rotation of the mirror is about an axis in the plane of the mirror, which is the case in an elevation scan, then $\mathbf{m}_0 \times \mathbf{m}$ coincides with that axis. If the angle between \mathbf{m}_0 and \mathbf{m} is α ($\mathbf{m}_0 \cdot \mathbf{m} = \cos\alpha$), then the rotation of the IFOV is 2α . We shall discuss separately the effect on the LOS and the effect on the IFOV function relative to the LOS.

In an elevation scan at zero azimuth, the LOS scans vertically at exactly twice the rotation rate of the mirror because the LOS is perpendicular to the elevation axis. However, in an elevation scan at nonzero azimuth, the LOS is not perpendicular to the elevation axis, so that it scans more slowly, and the direction of scan is inclined to the vertical. Both of these effects are very small (because the range of scan angles is so small), and it is only necessary to include them approximately in the LOS calibration algorithm (although the effects are, of course, exactly calculable).

There are two effects of scan mirror rotation on the IFOV function relative to the LOS. The first arises because the elevation–azimuth coordinate system on the sphere is nonuniform, so that when the IFOV function is rotated from one direction to another its gridded representation ought to change even though its shape is identical, because the coordinate grid at the new location is different from the old. We are able to ignore this effect because the range of elevation angles is so small (about $\pm 2^\circ$), which means that the IFOV function is confined to a narrow strip on the sphere where the grid is reasonably homogeneous. The other IFOV effect is image rotation. In the case discussed above of an elevation scan at nonzero azimuth, there is a component of the elevation axis parallel to the LOS, which induces a rotation of the image about the LOS. The effect of this is that the long axis of the detector element is no longer imaged horizontally into the atmosphere. We are able to ignore this effect because the range of *azimuth* angles is so small (about $\pm 4^\circ$), so that the rotation never exceeds about $7'$.

The effect on the instrument IFOV of rotation of the scan mirror can thus be deduced just from knowledge of the rotation involved. The rotation of the mirror about the two axes is measured by two linear variable differential transformers (LVDTs: Schaevitz 250MHR, Lucas Control System Products, Hampton, Virginia, and

Slough, United Kingdom) with improved electronics (designed by Rutherford Appleton Laboratory) which gave a noise level corresponding to 0.3" LOS motion in elevation. These telemeter the position of the mirror in both axes every IMP, encoding the result as a 16-bit count. The calibration requirements are much tighter for knowledge of the mirror position in elevation than in azimuth. The original design aims were as follows:

- Elevation: Relative line of sight direction between any two elevation scan positions not more than a quarter of the scan range apart shall be known to within 1.5" (three standard deviations).
- Azimuth: Relative line of sight shall be known to within 10" (three standard deviations).

These were set in the context of an estimated systematic error budget, with the intention that this source of error should not be the dominant instrumental error.

b. Data

The first attempt to calibrate the angle encoders was at the end of the first calibration campaign. The rotation of the scan mirror was directly measured with an autocollimating theodolite (Zeiss ETh3). For these measurements the instrument stood upside down on its mounting feet on the solid clean room floor. In this orientation the scan mirror normal is within a few degrees of the horizontal $+Y$ (or $-Y$) axis, and the mirror can be observed by opening one of the $\pm Y$ apertures in the outer structure. The autocollimation condition could be obtained over the whole scan range of the mirror, although there was some vignetting of the beam at the extreme "corners" of the scan range (i.e., at simultaneous extrema of azimuth and elevation). The dataset consisted of 80 theodolite settings together with the corresponding scan mirror step numbers and LVDT values averaged over the period of each setting.

Analysis of these data showed that the transducer angle/count characteristic is significantly nonlinear at the 10" level. Very similar nonlinearities were seen in both azimuth and elevation, and probably result from an intrinsic nonlinearity of the LVDT with distance rather than from the coupling to the angular motion or the electronics. The slope of the characteristic varies by about $\pm 1\%$ in azimuth over the scan range, and about $\pm 1.5\%$ in elevation. The dataset was clearly inadequate to characterize the nonlinearity adequately. A number of improvements were made for a second and much larger set of measurements at the end of the 1990 calibration campaign. Measurements were made in every elevation step on each of three scans at three different azimuth positions, followed by a single azimuth scan with an interval of five mirror steps, at the center of the elevation range; a total of 792 mirror positions over a 12-h period. The short-term repeatability was less than 1" (mirror rotation), but drifts in alignment between the

theodolite and the instrument occurred over longer timescales.

c. Transformation of measured angles into mirror reference frame

It is necessary to transform the measured angles from the theodolite reference frame to the frame defined by the mirror rotation axes. The transformation was established by using the mirror step numbers as approximate measures of the rotation about the mirror axes. Measurements at subsystem level had shown that the angle turned by the mirror was accurately described by a cubic polynomial in the step number and that the reproducibility of the steps was surprisingly high—about one-tenth of a step rms. The hysteresis between the average position for increasing and decreasing step number was also known from these measurements to be about 0.15 steps. The relationship between the two reference frames was determined by a least squares fit between the two angle sets: (i) the theodolite angles transformed into a different reference frame using three Euler angles and (ii) the angles defined by the step numbers (adjusted for the known hysteresis in the case of decreasing step numbers) using cubic polynomials in azimuth and elevation, with the six polynomial coefficients and three Euler angles as the unknowns.

The residuals in this fit were examined carefully for evidence of systematic effects. The overall rms deviation was 4.1", which is dominated by the random non-reproducibility in step sizes. There are, however, some systematic components in the residuals. Systematics in the scanning direction imply that the step-size variation is not purely random but contains a part that varies systematically with scan angle; these systematics are present, but not large in comparison with the random component. The more significant systematic residuals are in the nonscanning coordinate, such as the azimuth residual to a point in an elevation scan. Since there is nominally no motion in this coordinate, the random component here is much smaller, allowing the systematics to be more clearly seen. These imply that moving the mirror in elevation, say, can induce a small azimuth movement. These movements are not generally present in the LVDT counts and probably result from runout in the azimuth bearings. The possibility of mirror motion that is not correctly measured by the LVDTs clearly has significant consequences. The available evidence suggests that the unmeasured motion *in elevation* during an elevation scan is less than an arc second but that it would be undesirable to move the mirror in azimuth during an elevation scan. The operational scan pattern always makes azimuth movements at the top or bottom of the scan.

The fitted polynomial coefficients agree well with the independent step-size model derived from the subsystem data. The Euler angles are very well determined; for example, the angle between the azimuth axes of the two

reference frames is determined to be $12.60' \pm 0.10'$ (one standard deviation). Finally, this approach was extended to allow for the possibility that the two scan mirror axes are not orthogonal. No significant misalignment was found, with a limit of $0.6'$ (three standard deviations).

d. Calibration of LVDT counts

Having transformed the theodolite angles into the mirror reference frame, we have to address the nonlinear characteristic of the LVDTs. The solution that was adopted was based on the idea of a cubic spline, that is a continuous curve with locally constant first and second derivatives, but the variation in the third derivatives was "spread out" along the curve as required, instead of being concentrated at the spline knots. This was achieved with a Kalman smoother, in which the forward model extrapolates using a three-element state vector consisting of the local value and the first two derivatives of the angle/count characteristic, while the third derivative is assumed to be random with a certain variance. The fundamental assumption here is that the "true" angle/count characteristic is smooth and reproducible, but no other assumptions are made about it. The results of this fit were contained in a lookup table.

The smoother is implemented by combining two estimates of the true LOS angle θ at a given LVDT count L_i derived from Kalman filters running forward and backward through the data, which for this purpose consists of a single elevation scan. The degree of smoothing produced depends on the variance Q assigned to the stochastic third derivative. A constant third derivative q acting over a range of LVDT count ΔL generates a change in the θ of $q(\Delta L)^3/6$. We can therefore parameterize Q in terms of the distance ΔL over which the rms third derivative generates an additional uncertainty in the predicted curve equal to the measurement error. From the size of the nonlinear structure, ΔL is clearly on the order of 10 mirror steps, and varying it over a wide range gave the expected behavior of poor smoothing (following the data too well) if ΔL is too small and the emergence of systematic structure in the residuals (as a result of the smoothed curve not following the data) if ΔL is too large. Between these limits there is a wide range of values where the rms deviations between the smoothed values and the data vary very slowly with ΔL . The actual value used was 12 steps, which was selected as the value where the fit was least sensitive to ΔL . Having tied the value of Q to the measurement variance R , the value used for R does not affect the fit at all, only the statistical significance of the deviations between smoothed and measured values. Therefore, R was set so that the average squared difference was about R . The values used imply a measurement error of around $1.7''$ (mirror angle), significantly larger than the short-term repeatability, probably as a result of drifts in alignment over the time period required to take the approx-

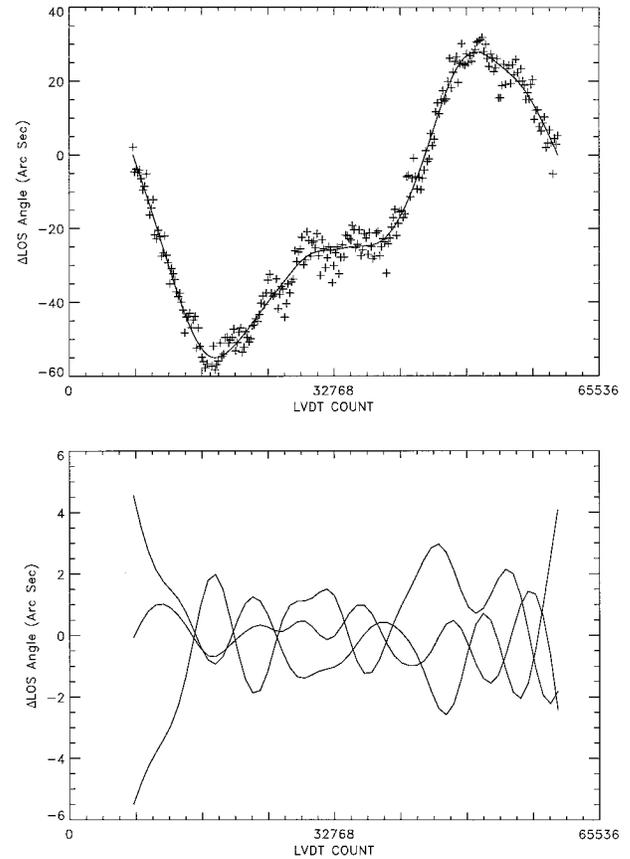


FIG. 4. Upper panel: Elevation scan angle encoder nonlinearity, shown as deviation from a straight line. Crosses: data, solid line: smoothed fit. Lower panel: Comparison of the three smoothed fits relative to their weighted mean.

imately 24 measurements contributing significantly to each smoothed value.

The output from the smoother consisted of 58 smoothed LOS angle values at predetermined and equally spaced LVDT counts for each elevation scan. One of the three elevation scans is shown in the upper panel of Fig. 4, as the deviation of the LVDT-angle characteristic from a straight line. The formal standard deviation of the smoothed values is about $0.8''$ over most of the scan. The three elevation scans were then combined to produce a single lookup table. The differences between the three scans significantly exceed their standard deviations ($\chi^2 = 3238$, 116 degrees of freedom). This is the result of longer term drifts in alignment between theodolite and scan mirror. There is no way of correcting these drifts within a given scan; we can only hope that averaging over the three scans reduces their effect. However, we can remove the average drift from one scan to the next by adding offsets to make the average drift zero. After this is done the differences are much smaller ($\chi^2 = 473$, 114 degrees of freedom). The final result is then taken to be the weighted mean of the three scans; the deviations from the weighted mean are shown in the

lower panel of Fig. 4. The best estimate of the error is taken to be the formal error in the weighted mean scaled by $(473/114)^{1/2}$; in the central portion of the scan range, this is very close to 1". Thus, we have not achieved our original aim, by about a factor of 2, but it is still the case that this error source is far from being the most significant overall (Livesey 1995; Dudhia and Livesey 1996).

The single azimuth scan was smoothed by a simpler procedure. Because the number of data points was comparable with the required number of points in the lookup table, instead of much greater in the case of the elevation scans, the lookup table was generated by simply smoothing the data. A simpler state vector was used, as the reduced information content is not sufficient to determine the second derivative. The accuracy of the smoothed values exceeds the requirement by a comfortable margin.

A small temperature correction was applied to these directly measured values. The normal operating temperature of the LVDTs in orbit is 11.5°C, which is several degrees above its passive equilibrium temperature, and maintained by a heater and thermostat. Because the angular calibration was done in ambient conditions, the LVDT temperature was about 10° warmer than this. The end-to-end temperature coefficient was measured after launch. This procedure took several hours, during which the mirror made repeated up and down elevation scans in single steps, while the set point temperature of the thermostat was increased in steps from its normal setting to its maximum. The analysis is based on the assumption that, on average, the mirror position in each step is independent of LVDT temperature, which is plausible as the drive and LVDT are at opposite sides of the mirror. The LVDT count recorded in each mirror step was regressed on the LVDT temperature to extract the temperature dependence; a linear temperature dependence and a linear dependence of the coefficients on mirror position were found to reproduce the data adequately. The effect is significant but not large: Over a quarter of the scan range, it amounts to 3", and the correction is estimated to be accurate to about 3%.

5. PMC calibration

a. The ISAMS pressure modulators

The ISAMS has eight PMCs, each containing a gas sample whose pressure is cycled at a particular frequency. The pressure modulator mechanisms are described more fully in Taylor et al. (1993). The pressure modulation is induced by a pair of opposed, oscillating pistons; the space between the pistons is connected to an optical cell through which is passed the radiation to be modulated. The effect of the gas in the cells is to selectively absorb the radiation emitted by the same gas present in the atmosphere, and the effect of the modulation is therefore to impose a time variation on this

component of the detected radiation. Each signal channel produces two outputs: one of which (the wideband output) is an average over this time variation, while the other (the pressure-modulated, or PM signal) is the result of phase-sensitive detection at the PMC frequency. Each of these signals is separately calibrated radiometrically, as described in Ballard et al. (1996), and has the same IFOV, because it is derived from the same detector signal. The significant difference between the two signals is therefore in their spectral response: The PM signal has greater selectivity than the wide band. This selectivity is illustrated in the case of NO₂ by Reburn et al. (1996), where the emission spectrum is heavily overlapped with H₂O emission, and for CH₄ and N₂O by Remedios et al. (1996), where both constituents emit in the same spectral region.

The instrument-level spectral response measurements are described in Ballard et al. (1996); in summary, the measurements are used to deduce the spectral response of the wideband channels in the absence of gas in the PMC, and the effect of the gas on the wideband and PM signals is then calculated using line-by-line computer techniques and standard spectroscopic line parameter compilations. This effort requires the mean gas state (temperature and pressure) for the wideband signal and the variation of these around the pressure cycle for the PM signal. As will be described below, it was not possible to make direct measurements of the varying component of the gas conditions in the flight modulators. Instead, these were determined from a computer model that was extensively validated against direct measurements using a development modulator. Calibration of the mechanisms thus consists largely of a determination of the mean gas pressure and temperature as a function of measurements available in the instrument telemetry stream. As this calibration must be determined in conjunction with precision pressure transducers unsuitable for flight use, most of this activity took place before installation of the mechanisms into the flight instrument.

b. Mean pressure

The mean pressure in the modulator is determined by the temperature of an amount of silicalite molecular sieve material, onto which is adsorbed the relevant gas. The sieve temperature, and hence the modulator mean pressure, can be set on command. The sieve material is contained in a separate unit, joined to the rest of the mechanism by a connecting pipe. The sieve unit is thermostated by a digitally controlled circuit; the thermostat setting, heater duty cycle, and sieve temperature are all available in the telemetry stream. Each sieve thermostat can be set independently over a range of approximately 0°–100°C.

The sieve temperatures were measured using precision rhodium-iron thermistors (nominally 27.3Ω at 273 K). There are uncertainties in this resistance measurement due to the inadvertent inclusion of about 1.2Ω of

contact resistance, variations in which resulted in temperature equivalent uncertainties of up to 2 K. In addition, the resolution of the analog-to-digital converter (ADC) channels used to measure the reference resistances against which the sensor resistances were calibrated was unnecessarily set to 0.1 K (about four times worse than that for the sensors). There was, in any case, a significant temperature drop across all of the sieve units (except channel 4), which were mounted on the exterior of the instrument to allow them to lose heat to cold space, which made interpretation of the measurements more difficult. The sieve temperature measurement, and any inferred pressure, can therefore only be regarded as approximate.

The sieve temperature measurements are used only as part of the leak measurements described in section 5e. Accurate measurements of modulator mean pressure are instead determined from the operating frequency of the modulator, making use of the variation of the resonant frequency with pressure. The gas between the pistons acts like a spring and so contributes a term to the resonant frequency; typically, a change of about 0.5 Hz mb⁻¹ is observed relative to a frequency when empty of about 30 Hz. The frequency of each PMC is measured in the instrument by comparison of the piston position transducer outputs with a 32-kHz signal derived from the spacecraft master clock, giving measurements with a resolution of better than 1 mHz over a 2-s interval. In practice, the limit on frequency knowledge appeared to be fluctuations of about 3 mHz rms, equivalent to about 0.006 mb, considerably worse than the accuracy and precision of the frequency measurement.

Calibration of the mechanism therefore consists largely of an accurate measurement of the pressure–frequency characteristic. The pressure was determined using one of two precision MKS Baratron 145-AH pressure sensors. These were not suitable for flight use so that the calibration had to be carried out before the modulator was finally filled, permanently sealed, and installed into the instrument.

Although the frequency depends primarily on the pressure, a number of other small effects must also be considered. The largest of these is caused by the variation with amplitude of the spring constant of the diaphragm springs supporting the modulator pistons. This causes the resonant frequency to vary with oscillation amplitude and also with the orientation of the piston axis when operated subject to a 1g gravitational acceleration. The latter effect results in a change of up to 0.4 Hz between vertical and horizontal operation. This had to be taken into account as much of the thermal-vacuum testing of the instrument was carried out with the axes vertical, whereas in orbit there is no 1g acceleration and the PMC performs as if horizontal.

In addition, the temperature of the modulator body, and the supply voltage to the output stage of the drive amplifier, also have small effects on the frequency for a given pressure; the frequency variations are typically

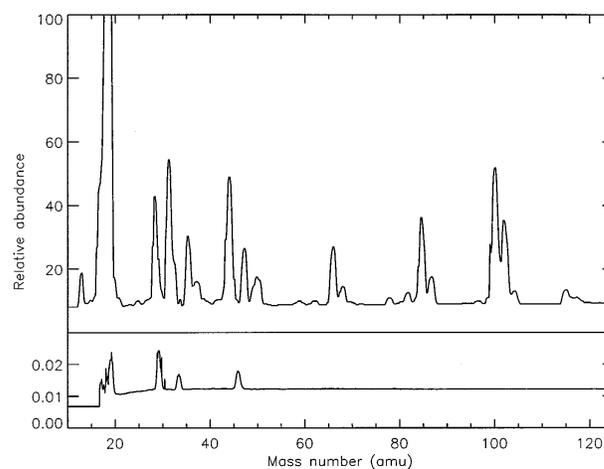


FIG. 5. Mass spectrum plots for the PMC filling rig. The upper plot shows measurements taken about 4 h after the beginning of pumping; the lower plot is after 478 h of pumping. Both plots have the same vertical scaling and an arbitrary zero level.

5 mHz K⁻¹ and 1 mHz V⁻¹, respectively. Although a voltage regulator was installed late in the program to eliminate the latter sensitivity, knowledge of the effect is still required to compare measurements taken at different times in the calibration program and in flight, when different voltages were used.

To summarize, measurement of the pressure–frequency characteristic required a thorough mapping of modulator operation against pressure, over a representative range of amplitudes, temperatures, and supply voltages. It is assumed that, once the modulator is sealed, this calibration remains valid throughout its life. The test program allowed a limited verification of this assumption, as two modulators (channels 0 and 1) were removed from the instrument and recalibrated. The repeatability of the calibrations, after instrument level vibration testing and 2000 h of operation, some 2×10^8 oscillations, was equivalent to 0.02 and 0.05 mb, respectively, for channels 0 and 1.

The calibration program involved thorough evacuation and baking of each of the modulators on a purpose built, chemically clean, filling rig. Each mechanism was evacuated for typically 10 weeks, at bake-out temperatures of 170°, 85°, and 65°C for the sieve unit, connecting pipework, and modulator body, respectively, at a typical base pressure of less than 5×10^{-7} mb. Isolated pressure rise measurements were made from time to time to check for leaks, although this measurement had only limited validity due to the necessary presence of an additional high vacuum pressure sensor (ion gauge), with a variable and comparatively high outgassing rate, in the isolated system. Further confidence in the cleanliness of the system was obtained by regular mass-spectrometer measurements; examples from early and late in the evacuation cycle for a typical modulator are shown in Fig. 5.

Once the evacuation was complete, and without re-

moving the modulator from the filling rig, the flight piston position transducer circuits and drive motor coils were fitted, and the pressure–frequency calibration carried out with the modulator axes horizontal. Throughout the calibration, pressure was measured by either of two Baratron gauges, one of 10 torr (13 mb) full scale and the other of 1000 torr (1300 mb) full scale. The latter unit was used only for the higher pressure (channels 2, 3, and 7) measurements. After all the modulators had been filled, the gauge calibrations were checked against traceable transfer pressure standards by the manufacturer.

The modulator operating frequency was measured by an HP5315A frequency meter, typically averaging over 30 s; modulator body temperature was measured with a Comark electronic thermometer. The operation of the PMC was controlled by an engineering model drive circuit of the same design as the flight boards; the drive amplifier duty cycle and piston position transducer outputs were also monitored.

The operating frequency was first measured as a function of amplitude demand (and at different body temperatures) with the modulator evacuated. The modulator was then filled with gas in stages, at each step allowing the pressure to equilibrate before again measuring the frequency–amplitude demand characteristic. For each modulator, the resultant dataset consisted of more than 100 combined pressure, frequency, and amplitude calibration points. A more restricted set of measurements of the small variations in operation with body temperature and drive voltage was made.

Once filled to the maximum required pressure and with the sieve at its nominal maximum temperature, typically about 100°C, the valves isolating the modulators from the filling rig were closed, and the sieve temperature cycled down to room temperature and up to the maximum again over about 15 h. This procedure was intended to fill all the molecular sieve sites with the gas. The mean pressure and the sieve temperature were then returned to their nominal maxima. When a further measurement of frequency against amplitude had been made, the copper filling pipe was double crimped off. For channels 4 and 5, which were filled with NO and NO₂, respectively, aluminium filling pipes were used. As a check, the frequency–amplitude measurements were then repeated, before the mechanisms were mounted onto handling fixtures prior to installation into the instrument.

A further test was carried out at instrument level to determine the variation in frequency between operation with the PMC axes horizontal and vertical. The former was the case for filling and (effectively) in-flight operations; the latter was the case for the instrument level calibration program. This test was carried out by rotating the whole instrument through 90° on its mount in the test chamber. As the effect was not expected to vary significantly with mean pressure, frequencies before and after rotation could simply be compared with the sieve

temperature held constant. The frequency shift was determined over the full range of amplitudes used for vertical axis operation. To verify the insensitivity to pressure, the tests were repeated over the reduced range of sieve temperatures accessible at instrument level in air (ambient to 50°C).

To determine the mean pressure in flight, functions were fitted to the calibration dataset giving pressure as a function of instrument telemetry; specifically, frequency, amplitude demand, body temperature, and supply voltage. Allowance was included for the different (7%) gains of the amplitude servo used in the calibration measurements compared with those on the flight instrument.

Theoretical considerations (Venters 1991) suggested that the operating frequency F should rise as the square root of the pressure P :

$$P = B(F^2 - F_0^2) = 2BF_0(F - F_0) + B(F - F_0)^2,$$

where F_0 is the operating frequency at zero pressure. In practice, a slightly more general function was used to provide a better fit to the data:

$$P = A(F_{\text{corr}} - F_{\text{empty}}) + B(F_{\text{corr}} - F_{\text{empty}})^2. \quad (1)$$

Here F_{empty} is the empty frequency as a function of PMC body temperature T :

$$F_{\text{empty}} = F_0 + C(T_0 - T), \quad (2)$$

where F_0 is the empty frequency at the standard temperature $T_0 = 25^\circ\text{C}$. In (1) F_{corr} is a corrected modulator frequency that takes account of the small variation with supply voltage V :

$$F_{\text{corr}} = F_{\text{obs}} + D(V_0 - V), \quad (3)$$

where F_{obs} is the observed operating frequency and V_0 the nominal supply voltage of 28 V. A necessary refinement, due to the variation in modulator effective spring constant with operating amplitude, was that the coefficients in the fitting functions were functions of amplitude. In all, a total of five amplitude-dependent coefficients were generated for each PMC: A , B , C , D , and F_0 . An additional two amplitude-dependent coefficients were also generated to fit the frequency change observed when the mechanisms were operated with axes vertical. A dependency not included in this pressure algorithm is a possible dependence of A and B on T . Some data on this are available, and improvements to the algorithm are under investigation.

c. Assessment of errors

Errors in the measurement of mean pressure can be divided into two classes: those arising from inaccurate modeling of the pressure–frequency characteristic by the fitting functions described above, and those arising from errors in the telemetered measurements. The former were estimated from the discrepancy between the predicted and measured pressures for the calibration mea-

surement data; the dominant errors in the calibration data were random, not systematic. The telemetered quantities required are PMC frequency, body temperature, drive amplitude, and supply voltage. The errors in the frequency measurement are discussed in section 5b.

The PMC body temperatures are derived from measurements of the temperature of the optics mounting plate (OMP), to which the mechanisms are attached. The temperature of the OMP is measured by precision-matched UUA41J1 thermistors at seven points and the temperature at each modulator mount then computed as a linear combination of these measurements. The water vapor modulator (channel 1) is insulated from the bench and the body temperature separately monitored and controlled. The thermistor resistances are measured using an accurate onboard 16-bit ADC, which is itself calibrated every minute in flight by making measurements of a precision reference resistance and of the appropriate ground reference. As the mean temperature is only a small correction to the pressure–frequency calibration, and has only a comparatively small effect on the spectral response of the instrument, the calibration of the thermistors and convertor were accepted as part of the PMC calibration without significant further instrument-level measurement. A simple check for gross errors was performed, however, by verifying that all sensors indicated the measured air temperature when the instrument was first powered up after being left overnight in a thermally stable environment; this test also allowed the variation of resistance offset errors in the sieve temperature measurements to be estimated.

Errors in the mean temperature measurements were dominated by uncertainty in extrapolation of the optics mounting plate temperatures to that of the PMC body. The temperature gradient across the OMP was typically 5°C. Making an allowance for power dissipation in the PMC itself, the uncertainty in the mean temperature was taken to be no worse than $\pm 1^\circ\text{C}$ under steady-state conditions, equivalent to an error in pressure of typically 0.01 mb.

The amplitude measurement uncertainty was dominated by errors in the transfer of measurements from the filling calibration data to the flight instrument; this was determined from direct measurement of the piston waveforms to be about 0.2% of amplitude, about one-half of the demand digitization level.

The small sensitivity of the pressure–frequency characteristic to the PMC drive voltage, which was discovered late in the program, results in the term D in Eq. (3). The voltage was measured to better than 0.1 V; uncertainties in this term were dominated by errors in modeling the sensitivity.

Table 2 lists the dominant error sources for each modulator as an equivalent pressure uncertainty. The total error has been found assuming that the individual terms are uncorrelated and can be root sum squared; errors have been computed for a representative range of operating amplitudes. Not shown are the “second order”

TABLE 2. Mean pressure uncertainties, millibar equivalent values.

| Channel | Temperature | Voltage | Amplitude | Algorithm | Total |
|---------|-------------|---------|-----------|-----------|-------|
| 0 | 0.0055 | 0.020 | 0.011 | 0.042 | 0.048 |
| 1 | 0.0100 | 0.032 | 0.008 | 0.048 | 0.058 |
| 2 | 0.0110 | 0.031 | 0.013 | 0.044 | 0.055 |
| 3 | 0.0119 | 0.045 | 0.012 | 0.113 | 0.122 |
| 4 | 0.0052 | 0.012 | 0.006 | 0.041 | 0.043 |
| 5 | 0.0078 | 0.017 | 0.010 | 0.034 | 0.039 |
| 6 | 0.0097 | 0.010 | 0.010 | 0.021 | 0.025 |
| 7 | 0.0116 | 0.155 | 0.002 | 0.078 | 0.174 |

error terms due, for example, to the variation of the voltage error term with temperature. These are estimated to be significantly smaller than the dominant terms listed.

d. Gas-state oscillations

Ideally, the pressure oscillations in the modulator optical cell would be measured directly and included in the instrument telemetry. However, at the time of design of the modulators (mid-1980s) no pressure transducers were available that met the performance, flight qualification, and chemical cleanliness requirements for IS-AMS. Accordingly, pressure cycles in the modulators were calculated using a computer model (Venters 1991), which was comprehensively validated using silicon diaphragm pressure sensors mounted in a development modulator.

The inputs to the model are the gas properties, mean gas pressure and temperature, and the modulator internal dimensions and piston amplitudes. The gas properties were determined from a number of standard sources (Venters 1991); the mean pressure and temperature were determined as described above. The modulator dimensions were taken from measurements of both individual parts and the assembled mechanisms, in conjunction with manufacturing drawings.

The final inputs required to the model were the piston amplitudes, which were derived directly from the instrument telemetry stream. This was possible because of the highly accurate (integrating) servo loop used to control the amplitudes in flight; both the amplitude measurement and the commanded setting are scaled relative to the same precision voltage reference. Under all normal conditions, the piston amplitudes are directly proportional to the digital commanded setting, provided only that the digital-to-analog convertor and the piston amplitude sensor gains are linear; this was verified at subsystem level.

Direct measurement of the temperature cycles is virtually impossible using contact sensors, due to the high frequency of the temperature oscillation in comparison with the thermal properties of any physically realistic sensor size. As it was also not feasible to include any radiometric measurements of the oscillation, the temperature cycle within the modulators was predicted by

the computer model. Confidence in this approach was based on extensive validation of the pressure cycle predictions of the model.

e. Gas composition

A concern with sealed pressure modulator-type instruments is the possibility of leaks. Over the period between sealing and launch, typically two years, the 100-to-1 pressure difference between the atmosphere and the PMC contents, combined with the comparatively small internal volume of about 1 L, can result in a significant accumulation of air in the modulator.

Once sealed, there is no measure of the pressure other than those described in section 5b. Assessment of a leak is made by comparing the apparent pressure, as determined from the frequency, amplitude, body temperature, and supply voltage telemetry, compared with what would be expected for a given sieve temperature. There is therefore no way of distinguishing between changes in the PMC pressure–frequency characteristic, changes in the sieve temperature measurement or pressure–temperature characteristic, and genuine leaks. To minimize possible ambiguity, great care must be taken with such measurements to ensure that identical conditions prevail over the extended period required to detect small leaks. This was very difficult to achieve in the range of thermal environments that the instrument encountered during prelaunch testing and was further hampered by the problems with the sieve temperature measurement discussed above.

Leak assessments have been made, with the assumption that the apparent leak is genuine. For at least two of the modulators (channels 2 and 6), the apparent leak continued when the instrument was under vacuum for periods of up to three weeks at a time during prelaunch testing. This has cast some doubt as to whether the leak is genuine. It is hoped that data obtained from the much longer period of operation in vacuum after launch will enable the situation to be clarified; this is the subject of continuing investigations.

6. Conclusions

In Part I (Ballard et al. 1996) and this paper we have discussed a number of design errors and calibration problems with the ISAMS instrument; in particular, the scan stray radiance, the apparent inconsistency between the spectral and radiometric calibrations, the multiple reflections within the detector package, and the uncertainty in some PMC contents. Of these, the effects of the multiple reflections are fully characterized and are not major sources of error. We shall briefly refer to discussions of the remaining problems in the various channels.

The situation for the filter wheel channels is simplest, in that there is no significant inconsistency between the spectral and radiometric calibrations, and in the case of

the temperature-sounding channels there is no evidence of a significant PMC leak. On the other hand, the scan stray is significantly worse in the filter wheel channels. Livesey (1995) has discussed all the error sources relevant to the temperature-sounding channels in some detail, and Dudhia and Livesey (1996) summarize the conclusions. Connor et al. (1996) and Smith et al. (1996) give error estimates for the O₃ and N₂O₅ filter wheel channels.

In the remaining channels, all these problems are present to a greater or lesser extent, and the error budget is discussed by Goss-Custard et al. (1996) (channel 1, H₂O), Remedios et al. (1996) (channels 2 and 6, N₂O and CH₄), and Reburn et al. (1996) (channel 5, NO₂). In addition to these instrumental problems, the retrievals are significantly complicated by the unusually large atmospheric aerosol loading during the mission; the aerosol measurements are discussed by Lambert et al. (1996). The largest apparent spectral shifts are in channels 0 and 4. The situation in channel 4 is further complicated by the observation of significant emission from thermospheric NO, and there are at present no publicly available data from this channel. The interpretation of the channel 0 radiances requires the consideration of processes not in local thermodynamic equilibrium, and the error analysis is discussed by Lopez-Valverde et al. (1996).

Finally, we wish to draw some general conclusions based on our experience with ISAMS calibration. Calibration has to be done at the end of an instrument program, so that unexpected effects are only discovered at a very late stage; certainly some of the effects on ISAMS could have been eliminated had they been discovered earlier. This is a strong argument for making representative measurements on an engineering model instrument; in the case of ISAMS, there was a radiometric test model, with fewer channels and simplified optics, but its test program was overtaken by that of the flight instrument, so that certain measurements on it were never made. Program delays or budget constraints mean that this is often the case.

One aspect of our calibration approach, as used on ISAMS and the Along Track Scanning Radiometer, is the use of panels mounted closely around the instrument to simulate the radiative boundary conditions on each face of the instrument. This results in a close match between many instrumental temperatures in the test and orbital environments, and we consider this an important consideration for calibration of a midinfrared radiometer.

Control of the instrument and the test equipment from one node (the ISAMS microprocessor system), and the ability of the test team to program tests in a simple high-level language, made performing complex tests under great time pressure very much more straightforward. Similarly, the availability of the instrument and test equipment telemetry in a single coherent database greatly simplified subsequent data analysis.

The involvement of the principal investigator's institution in the calibration has been a great benefit: Satellite instruments tend to have unexpected features that become well known to those who command the instrument during prelaunch test and calibration, and the sharing of that experience, and to some extent the personnel, between the calibration, flight operations, and data analysis teams is very valuable. We note that this model is being followed for the measurement of pollution in the troposphere instrument on the first NASA Earth Observing System platform and for the high-resolution dynamics limb sounder instrument on the EOS Chemistry platform.

Acknowledgments. We wish to acknowledge the involvement of C. J. Scheuer of the Science and Technology Corporation (Hampton, Virginia) in the data analysis for the IFOV map.

REFERENCES

- Ballard, J., and Coauthors, 1996: Calibration of the Improved Stratospheric and Mesospheric Sounder. Part I: Introduction, spectral, and radiometric calibration. *J. Atmos. Oceanic Technol.*, **13**, 810–831.
- Connor, B. J., C. J. Scheuer, D. A. Chu, J. J. Remedios, R. G. Grainger, C. D. Rodgers, and F. W. Taylor, 1996: Ozone in the middle atmosphere as measured by the Improved Stratospheric and Mesospheric Sounder. *J. Geophys. Res.*, **101**, 9831–9841.
- Dudhia, A., and N. J. Livesey, 1996: Validation of temperature measurements from the Improved Stratospheric and Mesospheric Sounder. *J. Geophys. Res.*, **101**, 9795–9809.
- Goss-Custard, M., and Coauthors, 1996: Measurements of water vapour distributions from the Improved Stratospheric and Mesospheric Sounder: Retrieval and validation. *J. Geophys. Res.*, **101**, 9907–9928.
- Lambert, A., and Coauthors, 1996: Validation of aerosol measurements from the Improved Stratospheric and Mesospheric Sounder. *J. Geophys. Res.*, **101**, 9811–9830.
- Livesey, N., 1995: Small scale features in atmospheric temperature fields measured by ISAMS. Ph.D. thesis, University of Oxford, 235 pp.
- Lopez-Valverde, M. A., M. Lopez-Puertas, J. J. Remedios, C. D. Rodgers, F. W. Taylor, E. C. Zipf, and P. W. Erdman, 1996: Validation of measurements of carbon monoxide from the Improved Stratospheric and Mesospheric Sounder. *J. Geophys. Res.*, **101**, 9929–9955.
- Reburn, W. J., and Coauthors, 1996: Validation of nitrogen dioxide measurements from the Improved Stratospheric and Mesospheric Sounder. *J. Geophys. Res.*, **101**, 9873–9895.
- Remedios, J. J., and Coauthors, 1996: Measurements of methane and nitrous oxide distributions by the Improved Stratospheric and Mesospheric Sounder: Retrieval and validation. *J. Geophys. Res.*, **101**, 9843–9871.
- Rodgers, C. D., R. J. Wells, R. G. Grainger, and F. W. Taylor, 1996: Improved Stratospheric and Mesospheric Sounder validation: General approach and in-flight radiometric calibration. *J. Geophys. Res.*, **101**, 9775–9793.
- Smith, S. E., and Coauthors, 1996: Dinitrogen pentoxide measurements from the Improved Stratospheric and Mesospheric Sounder: Validation of preliminary results. *J. Geophys. Res.*, **101**, 9897–9906.
- Taylor, F. W., and Coauthors, 1993: Remote sensing of atmospheric structure and composition by pressure modulator radiometry from space: The ISAMS experiment on UARS. *J. Geophys. Res.*, **98**, 10 799–10 814.
- Venters, P. 1991: The pressure modulation system in the Improved Stratospheric and Mesospheric Sounder. Ph.D. thesis, University of Oxford, 249 pp.